ESD-TR-77-228

ESD ACCESSION LIST
DRI Call No. 87633
Copy No. 1 of 2 cys.

Technical Note 1977-37

P. Demko

Measurements of Voice and Data Queue Behavior in a PVC Network Link

25 August 1977

Prepared for the Defense Communications Agency
under Electronic Systems Division Contract F19628-76-C-0002 by

Lincoln Laboratory

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

LEXINGTON, MASSACHUSETTS

ADA047101

Approved for public release; distribution unlimited.

The work reported in this document was performed at Lincoln Laboratory, a center for research operated by Massachusetts Institute of Technology, for the Military Satellite Office of the Defense Communications Agency under Air Force Contract F19628-76-C-0002.

This report may be reproduced to satisfy needs of U. S. Government agencies.

The views and conclusions contained in this document are those of the contractor and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the United States Government.

This technical report has been reviewed and is approved for publication.

FOR THE COMMANDER

Raymond L. Loiselle

Raymond L. Loiselle, Lt. Col., USAF
Chief, ESD Lincoln Laboratory Project Office

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

LINCOLN LABORATORY

MEASUREMENTS OF VOICE AND DATA QUEUE BEHAVIOR IN A PVC NETWORK LINK

P. DEMKO

Group 24

TECHNICAL NOTE 1977-37

25 AUGUST 1977

Approved for public release; distribution unlimited.

LEXINGTON MASSACHUSETTS

ABSTRACT

This note describes measurements made on a computer simulation of a model of a Packetized Virtual Circuit (PVC) network link. The simulation models a population of speakers in conversation, and a Poisson data source. Such variables as buffer space requirements, packet loss and delay, and link utilization are investigated as functions of voice and data loads on the system. Initial results indicate that voice and data can be satisfactorily integrated on a 1.544 Mbps packetized communications link with total link utilizations greater than 90%.

TABLE OF CONTENTS

I. INTRODUCTION

Packetized Virtual Circuit (PVC) techniques[1] combine features of both circuit and packet switching to provide an efficient approach to integrating voice and data in a communications network. The PVC approach handles both voice and data traffic in an essentially uniform fashion, easing the implementation and providing the capability to respond automatically to changes in traffic mix. The PVC network concept, while attempting to capitalize on the statistical multiplexing advantages inherent in packet technology, also attempts to overcome some of the efficiency and delay dispersion difficulties associated with pure packet networks by utilizing communication link formats and routing conventions associated with digital circuit switching.

In the PVC network, connections are established from source to destination hosts. Since all packets in the connection follow the same path through the network, the PVC packet header need only contain information identifying it as belonging to a particular connection and packet overhead is reduced significantly.

In the PVC scheme a relatively efficient "statistical flow control" is performed by controlling the assignment of connections to specific links such that the probability of internal overloads is reduced.

A model of a single link of a PVC network has been developed and simulated[2] on a PDP-11/45 computer. The simulation models a population of speakers in conversation, providing a voice load on the system. Data traffic is modeled by a Poisson process. The PVC link model permits the investigation of such variables as buffer space requirements, packet loss and delay, and link utilizations as functions of the voice and data loads on the system.

Some characterisitcs of traffic consisting of packetized, encoded voice with silence detection were measured. Priority schemes and queueing strategies were developed to utilize most of the remaining channel capacity (not used by voice) by data traffic in a manner such that the voice traffic is not significantly disrupted and the data packet queue and resulting delays are not excessive. Initial results indicate that voice and data can be satisfactorily integrated with total link utilizations greater than 90%.

II. THE PVC LINK MODEL

A PVC link is modeled as having two distinct input queues, one for voice and the other for data traffic, as shown in Figure 1. When the link is available, a packet is chosen from one queue or the other and transmitted. Nemeth[2] describes the details of the model; a summary of the fixed parameters of the model is presented in Table 1.

For each run of the simulation, the number of speakers using each vocoder type is specified. Each speaker is determined to be speaking (active) or silent according to distributions of talkspurt and silence distributions obtained from measurements by Brady[3]. When a speaker is determined to be active, he generates packets at a rate characteristic of his vocoding technique. When he is silent, no packets are generated. The model does not attempt to represent the start or end of conversations. When a voice or data packet is generated, it is entered into the respective queue for transmission. The voice queue is finite; when it is filled, half the packets in the queue are discarded. The maximum size of the data queue is a variable and is measured for different traffic loads.

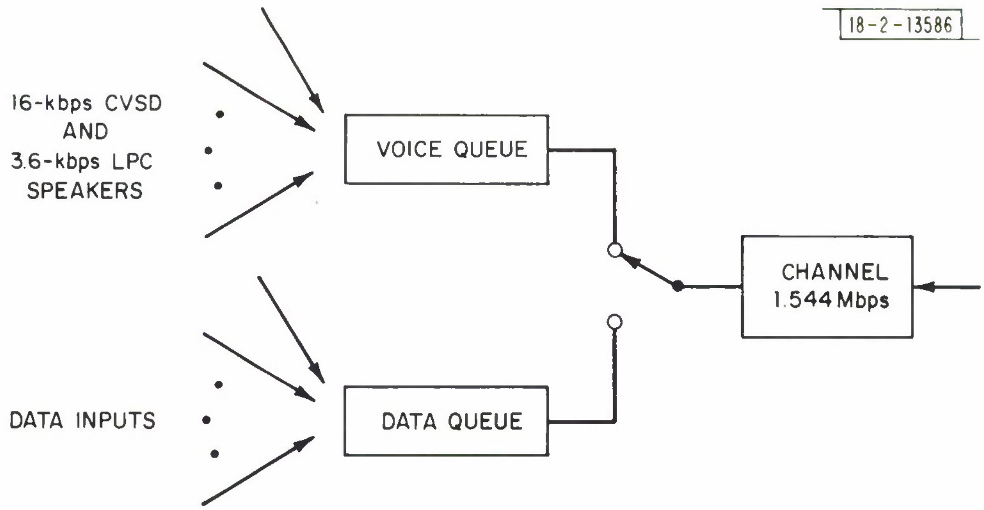

Fig. 1. Model of a single link in a PVC network.

TABLE 1

FIXED PARAMETERS IN THE LINK MODEL

Packet size	128 bits
Overhead in packet	32 bits
Data in packet	96 bits
Channel rate	1.544 Mbps
Supervisory traffic and framing	8 kbps
Available channel rate	1.536 Mbps 12,000 packets/sec
Vocoding techniques	CVSD, LPC
CVSD vocoding rate	16 kbps 6 msec between packets
LPC vocoding rate	3.5 kbps 27.5 msec between packets
Voice queue size	70 packets 560 16-bit words 5.83 msec of channel time
Simulation duration	2 min. of channel time

III. BEHAVIOR OF VOICE AND DATA TRAFFIC IN THE PVC LINK SIMULATION

A. Voice Queue

Initial measurements on the simulation were made with only voice traffic. It was assumed that voice had absolute priority on the link; measurements were collected on the remaining, unused portion of the channel to predict the behavior or the data channel (see next section). Nemeth's report[2] describes the results from these simulation runs. Various parameters were measured with a different total number of speakers for each run: the utilization of the channel capacity for voice traffic, the mean duration of contiguous voice slots (packet times) or mean time between empty packet slots, and the rate of data traffic that could be transmitted over the channel if all empty (non-voice) slots were used for data.

Loss of voice packets due to overflow of the voice queue was observed to occur only in those cases where the channel utilization for voice packets, ρ_v, was greater than 80%. Thus, theoretically, 20% of the channel or 230 Kbps is available for data traffic. The remainder of this report addresses the issue of how much of this capacity can be used for data traffic with appropriate storage and delay constraints.

B. Data Queue

1. Queueing Theory Approach

Given the statistics of the unused fraction of the channel measured above, one can attempt to predict the behavior of the data queue by modeling the data queue as an M/G/1 queue[2]. The gaps between empty slots can be considered as service times for data packets. The use of the M/G/1 queue model requires the following assumptions:

1) The data packets arrive according to a Poisson model

2) Voice traffic has absolute priority

3) Successive service times (gaps between non-voice slots) are independent.

The first two assumptions are clearly met in the model; the third assumption is questionable and will be discussed below.

The M/G/1 model provides [4] a formula for the mean waiting time, $\overline{W}$, for a packet in the data queue as a function of data arrival rate. The formula can be used to estimate link utilization as a function of the desired mean waiting time for data packet transmission. For example, calculations indicate that if one chooses a population of 100 CVSD speakers, a load just slightly greater (104%) than the trunk could handle with pure circuit-switching, the predicted mean wait is 8.3 milliseconds when transmitting 320 Kbits/second of user data (427 Kbits/second when packet overhead is included). Under these conditions no voice packets are lost and the worst case delay for voice is less than 6 milliseconds. Voice traffic is using 70.6% of the packet slots and the queueing model predicts that data occupies 94.5% of the remaining slots. Total channel utilization is 98.4%. If packet overhead is considered, the net utilization for voice and data is 73.8% of link capacity.

The suspect assumption for using the M/G/1 queue model - independence of successive service times - was checked by computing the first serial correlation coefficient of the duration of successive contiguous voice packet intervals. Following Cox and Lewis[5], $\tilde{\rho}_1$, is the unbiased estimate of the first serial correlation coefficient:

$$\tilde{\rho}_1 = \frac{\sum_{i=1}^{n-1} \left[x_i - \frac{1}{n-1}\sum_{i=1}^{n-1} x_i\right]\left[x_{i+1} - \frac{1}{n-1}\sum_{i=1}^{n-1} x_i\right]}{\sum_{i=1}^{n-1} \left[x_i - \frac{1}{n-1}\sum_{i=1}^{n-1} x_i\right]^2} \tag{1}$$

The quantity, $\tilde{\rho}_1 \sqrt{n-1}$ will have a unit normal distribution if ρ_1, the actual correlation coefficient, is zero and n is large. Independence is rejected as a hypothesis at the α significance level if

$$|\tilde{\rho}| > \frac{C_{1/2\alpha}}{n-1} \quad , \tag{2}$$

where $C_{1/2\alpha}$ is the upper $1/2\alpha$ point of the unit normal distribution.

The correlation measurements are summarized in Table 2. The successive intervals between empty slots fail this initial test for independence at significance levels of 5%, 2%, and 1%. It follows then that the M/G/1 queue model may not provide accurate predictions for the behavior of the data queue and the performance figures for the channel utilization quoted above are probably not reliable. Since queueing theory offers no mechanism for analytically coping with correlated service times, an augmented simulation that included data traffic was deemed to be appropriate. The queueing theory predictions and simulation measurements are briefly compared in the next section.

TABLE 2

COMPUTATIONS OF FIRST SERIAL CORRELATION COEFFICIENT
100 CVSD SPEAKERS
($n = 65,536$)

Estimates of Coefficient $\tilde{\rho}_1$
−0.051759
−0.054166
−0.023981
−0.036636
−0.033078
−0.013363

Significance (α) Level (percent)	$C_{1/2\alpha}/\sqrt{n-1}$
5	0.007656
2	0.009102
1	0.010060

2. Experimental Measurements on the Data Queue

The basic measurement for determining the behavior of data traffic is a histogram indicating the number of occurrences of each length (number of packets) of the data queue. The average waiting can be computed using Little's result[6].

$$\overline{W} = N_q/\lambda, \tag{3}$$

where N_q is the average number of packets in the queue and λ is the average arrival rate of packets per slot time. N_q can be calculated from Eqn. 4,

$$N_q = \frac{\Sigma i.n_i}{N}, \tag{4}$$

where i is the length of the data queue in packets, n_i is the number of slot times which that length queue occurred during the simulation run, and N is the total number of slot times during the run. Recognizing that λ is the average number of data packets transmitted during the simulation run, the mean waiting time becomes

$$\overline{W} = \frac{\Sigma i.n_i}{n_d}, \tag{5}$$

where n_d is the total number of data packets transmitted. The measured waiting time* is compared to that predicted by the M/G/1 queue in Figure 2.

---

*When the PVC link is run, a software random number generator produces the sequence of numbers used to select talkspurt and silence durations from their respective distributions. Usually, the number generator is started with a random "seed" for each run of the simulation. However, when computing data points from many simulations for a particular curve or family of curves, the same "seed" is used. The specific effects of one variable (for example, input rate of data traffic) can thus be analyzed without the dispersion in results due to run-to-run variations in speaker activity.

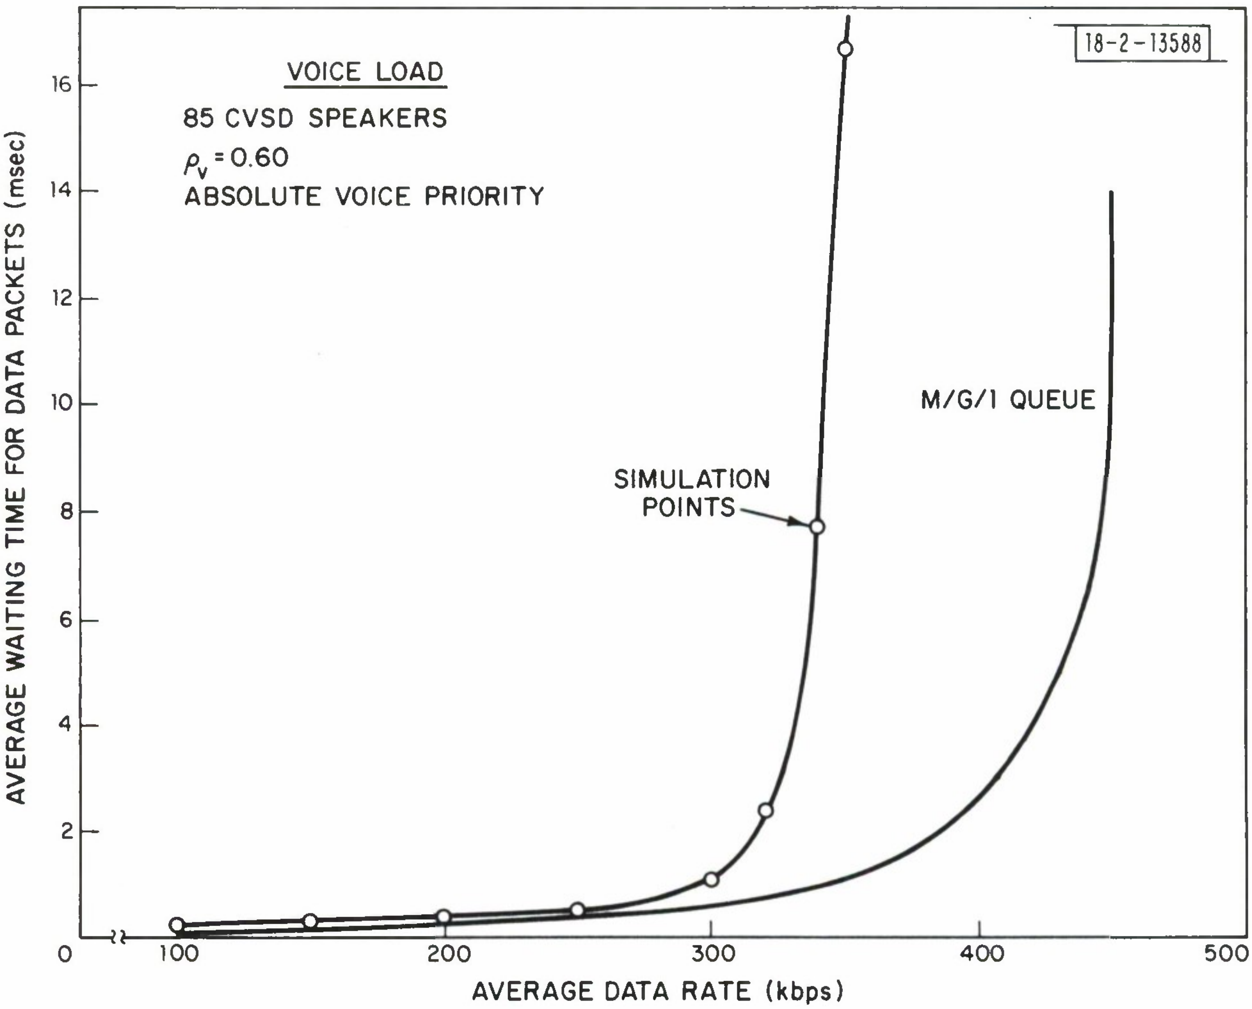

Fig. 2. Average waiting time vs. data rate.

The results from the simulation indicate that the fraction of the channel not used by voice traffic cannot carry as much data at the same delay as that predicted by the M/G/1 queue. For a 6 millisecond delay, the M/G/1 queue predicts a data load of approximately 450 kbps and total link utilization of 99%, while the simulation measurements indicate a data load of about 330 kpbs and total link utilization of 90%.

In summary, the measurements indicate that although, in addition to voice traffic, there is sufficient net capacity available for a rate of data traffic that brings the total link utilization to 98-99%, the statistics of voice traffic with absolute priority are such that very large delays in data packet transmission and unacceptably large queue lengths result. Nonetheless, one can maintain acceptable delays and queue lengths when data traffic is introduced at rates that result in net link utilizations of 90-92% - still relatively high values.

IV. ADDITIONAL MEASUREMENTS ON A SINGLE PVC LINK

A. Varying Data and Voice Priorities

When data traffic was introduced into the simulation, the display routines were modified to display the behavior of the data queue as well as the voice queue. When voice traffic has absolute priority, data packets are transmitted only when the voice queue is empty. It is observed on the display that often the voice queue remains small, but non-zero, for extended periods of time; thus, data transmission is delayed and despite the fact that the average data rate is close to anticipated values, a very large data queue results. Consequently, different priority strategies for transmitting voice and data were investigated.

A framing strategy can be introduced in which a fraction of the packet slots in the frame have priority for data. The fraction can be made to vary according to the voice and data loads in the node.

Figure 3 presents the same data as Figure 2 with the addition of measurements made when voice packets had priority for only 7 out of every 10 packets. The change in priority decreases the mean wait for data and the maximum size of the data queue and increases somewhat the delay for voice packets. For the voice load depicted in Figure 3, the decrease in the number of packet slots with voice priority did not increase voice packet delay significantly enough to result in speech loss. However, in cases with larger voice loads which result in no speech loss with absolute voice priority, there is speech lost when some priority is given to data.

A more detailed picture of the effects of changing voice and data packet priorities is shown in Figure 4. The data in Figures 4a, b, c, and d are all plotted against the number of packets (out of 10) with voice priority. The data and voice load on the link is exceptionally heavy and would not be used in a practical situation. Such a traffic load provides a large dynamic range of speech loss and data waiting time such that the effects of varying the voice/data priorities may be observed. The speaker load alone can potentially utilize about 80% of the link capacity while the data load alone can potentially utilize about 30%. When voice has zero priority, voice packets are transmitted only when the data queue is empty and approximately 11% of the speech is lost (when the voice queue overflows, half its packets are discarded). This speech loss does not decrease substantially until 80% of packet slots have priority for

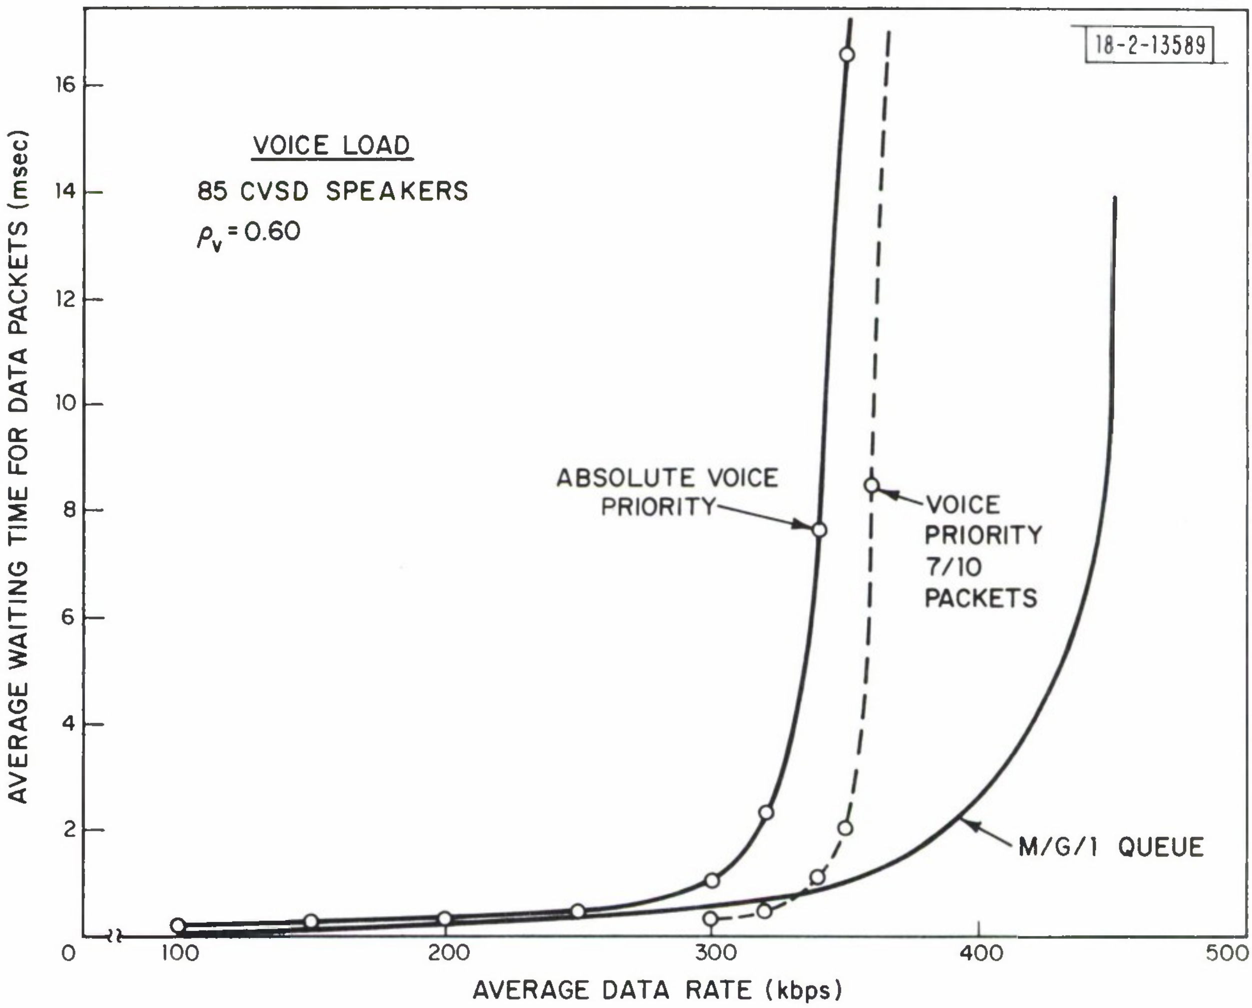

Fig. 3. Average waiting time vs. data rate.

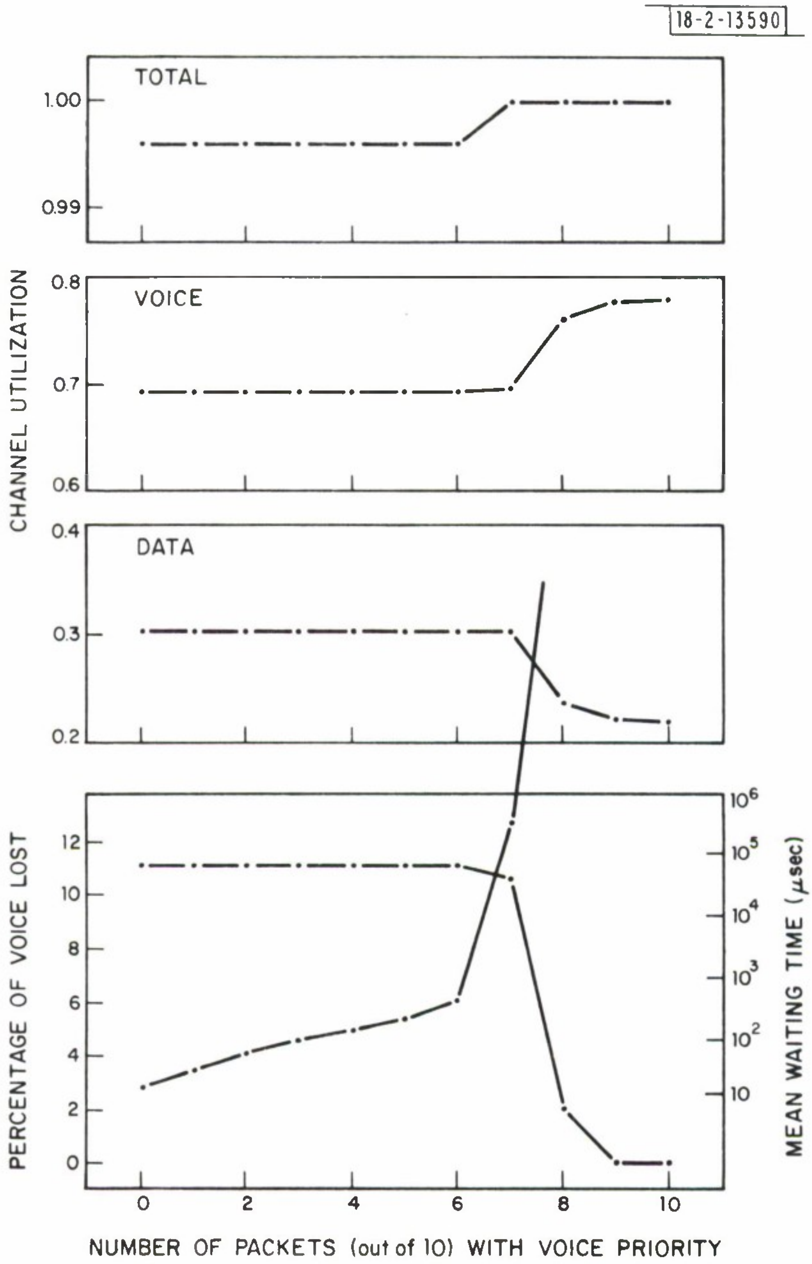

Fig. 4. Voice and data priority.

speech. When speech has all or nearly all the priority, it utilizes nearly 80% of the link capacity without any queue overflow. Data packets fill in the remaining packet slots resulting in a net link utilization of almost 100%. However, since more data packets are presented to the link than can be handled, queued data packets build up indefinitely.

B. Voice Queue Overflow Strategy

When the simulation was initially constructed, it was decided to discard half the queue (35 packets) when there was an overflow. It was thought that discarding the packets in this manner might spread the packet loss more evenly over all of the active speakers. An appropriate question is, for a heavy traffic load, can the amount of speech lost be decreased if fewer packets are discarded during overflows and/or if the size of the voice queue is increased (increasing the maximum delay for speech packets). The answer to the question is shown in Figure 5. For a given set of voice and data traffic loads and priorities the percentage of speech lost is plotted against the size of the voice queue. There are two curves: one in which half the queue is discarded at overflows, and the other in which only incoming packets are discarded when the queue is full. Several results can be discerned from the graph. Speech loss cannot be significantly decreased by increasing the size of the voice queue. The increase in maximum packet delay probably outweighs the small decrease in loss. It also appears that voice queue overflows occur in bursts, since discarding incoming packets only does not save appreciably more speech than discarding 35 packets when the queue reaches capacity. The optimal queue size is then between 50 and 100 packets (4.167 to 8.33 milliseconds maximum delay) and only incoming packets ought to be discarded when the voice queue is full. Alternately, if discarding half the queue spreads the packet loss more evenly over all the speakers (this fact is currently

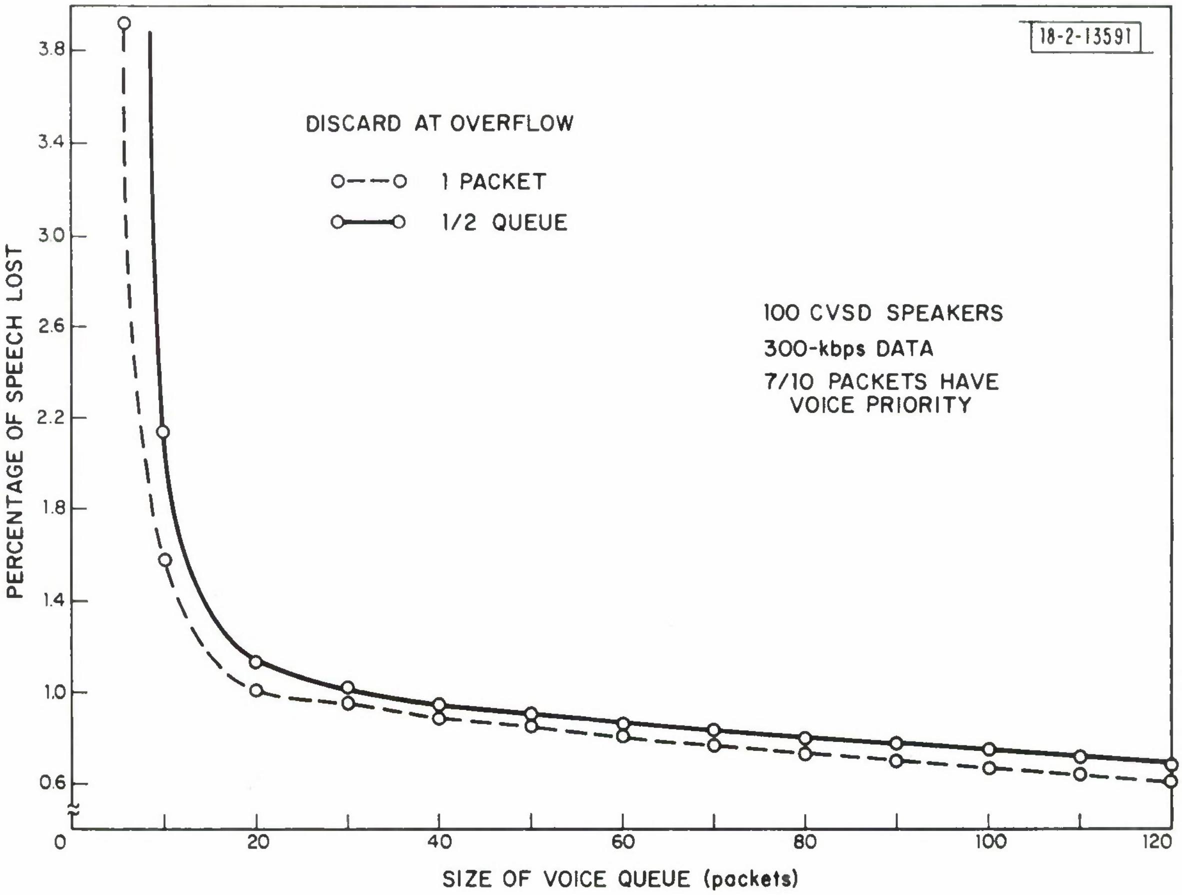

Fig. 5. Speech loss vs. size of voice queue.

not known), the small increase in total speech loss would be justified.

C. Behavior of Smaller Capacity Links

The original PVC link simulation assumes a link capacity of 1.544 Mbps. When speech activity detectors are used and no voice packets are transmitted during silence, a 1.544 Mbps link can handle approximately 100 to 125 16Kbps CVSD speakers with minimal speech loss. Experience[7] shows that the TASI advantage can safely be used only when the capacity of the channel shared by the conversations is relatively large (the order of 50 to 100 conversants). How does the PVC concept fare in networks with smaller capacity links? More specifically, given the smaller capacity links that cannot benefit from the TASI advantage, can the remaining capacity (unused by voice) be used (with acceptable delays) by data? These questions were investigated and the results follow.

In Figure 6 the percentage of speech lost is plotted against the utilization of the channel for voice for several link capacities. No data was transmitted. Clearly, as is predicted in the literature[7], smaller capacity links cannot use the TASI advantage as well as those with larger capacity. At a speech loss level of 1%, 93% of a 1.544 Mbps channel is utilized for voice traffic, while only 73% of a 128.64 Kbps channel is utilized for voice traffic.

A 0.1% speech loss level was selected and voice loads were determined from the curves in Figure 6 for the 128.640 Kbps and 1.544 Mbps links. For the smaller-capacity link, 7 speakers result in a link utilization of 0.59; for the larger-capacity link, 123 speakers result in a link utilization of 0.87.

The PVC link simulation was run at these voice loads with varying data loads; the data rates were restricted by requiring acceptable mean waiting times and queue lengths. The utilization of the link by data and the mean waiting time

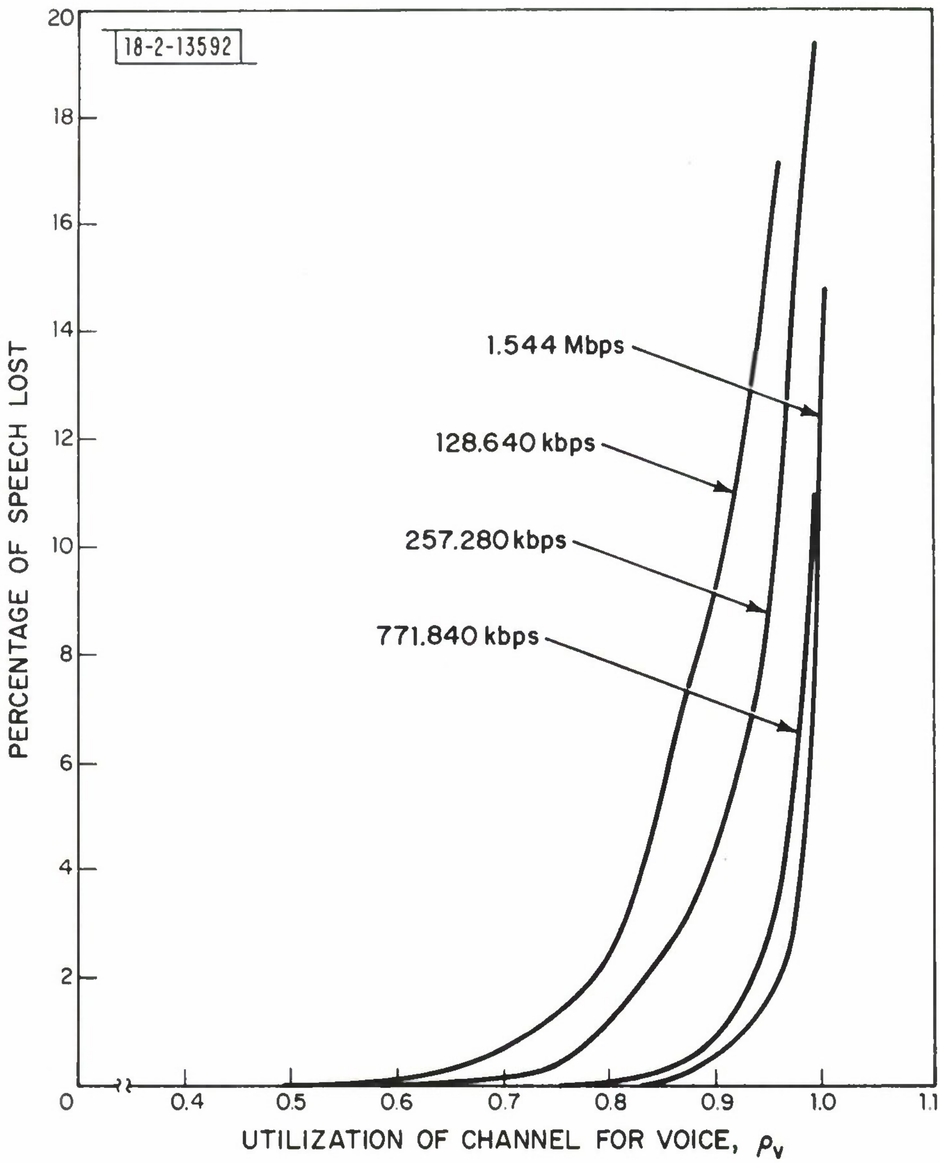

Fig. 6. Speech lost vs. utilization for different link capacities.

for data packets were measured. The results are plotted in Figure 7. Link utilizations, ρ 's, are plotted on a linear scale; mean waiting times, $\overline{W}$'s, are plotted on a log scale. Simulations were first run with absolute priority of voice over data. For the 128.640 Kbps link, utilization of the link for data ranged from 0.052 to 0.312 with mean waiting times ranging from 54.4 msec to 1027.1 msec. The maximum length of the data queue varied from 366 to 3434 data packets. For the 1.544 Mbps link, the data utilization ranged from 0.022 to 0.087 with mean waiting times ranging from 35.8 msec to 641.7 msec. The length of the data queue varied from 310 to 4834 data packets.

The utilization of the link by data in the smaller link is 2 to 4 times that of the larger capacity link, but the net utilization (including packet overhead) is still only 0.689 while that of the larger link is 0.913.

Further simulations were run with some priority given to data packets. In the 128.640 Kbps link voice was only given priority for 70% of the slots (ρ_V=0.59); in the 1.544 Mbps link, voice was given priority for 90% of the slots (ρ_V=0.87). The results of these runs are shown with dashed curves in Figure 7. The changes in priority increases speech loss to 5.0% in the smaller capacity link and 1.4% in the larger. Mean waiting time decreases significantly, but utilization of the links by data does not increase.

One can conclude that the high link utilizations by voice and data which result from PVC techniques are not achieved with significantly lower capacity links. Although the lower capacity link can carry proportionally more data, the utilization of the link by voice dominates the overall channel utilization for the given transmission priorities.

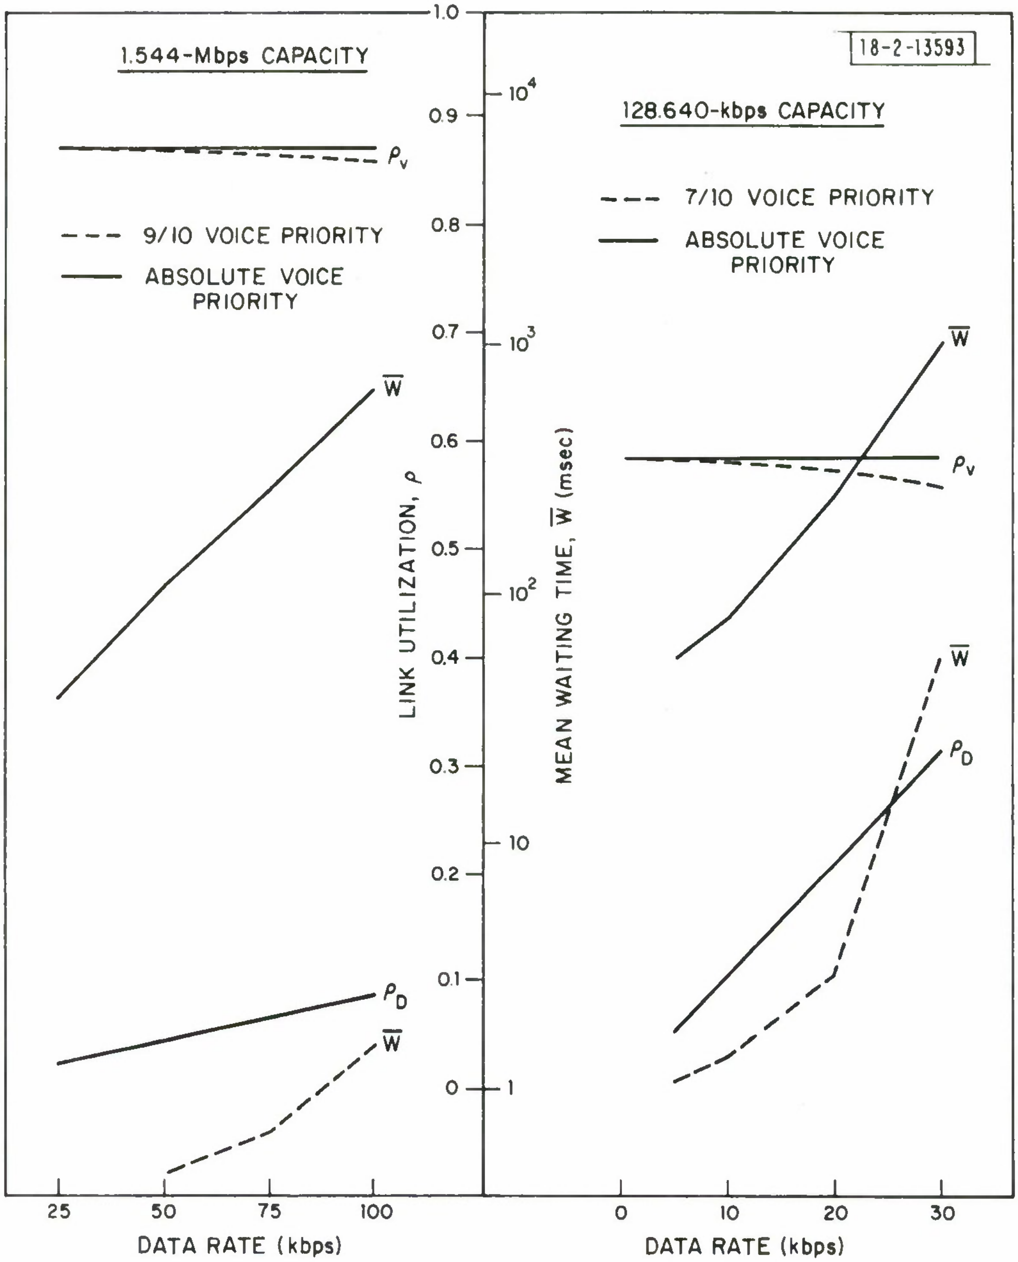

Fig. 7. Link utilizations for voice and data (linear scale). Mean waiting times for data packets (logarithmic scale). Solid curves: absolute voice priority. Dashed curves: indicated voice/data priority (see text).

ACKNOWLEDGEMENTS

The original concept for the PVC network was developed by James Forgie. Alan Nemeth wrote the initial model simulation of the PVC link. Clifford Weinstein provided the author with invaluable suggestions during the course of the work.

REFERENCES

1. Forgie, J.W. and A.G. Nemeth, "An Efficient Packetized Voice/Data Network Using Statistical Flow Control," Proc. International Communications Conference, ICC 77, Vol. III, pp. 44-48 (June 1977).

2. Nemeth, A.G., "Behavior of a Link in a PVC Network," Technical Note 1976-45, Lincoln Laboratory, M.I.T. (7 December 1976), DDC, AD-A036370/5.

3. Brady, P.T., "A Technique for Investigating On-Off Patterns of Speech," Bell System Technical Journal, 44, (January 1965), pp. 1-27.

4. Kleinrock L., Queueing Systems, Vol. 2: Computer Applications, (Wiley, New York, 1976), 16-19.

5. Cox, D.R. and P.A.W. Lewis, The Statistical Analysis of Series of Events, (Methuen, London, 1966).

6. Kleinrock, L., Queueing Systems, Vol. 1: Theory, (Wiley, New York, 1975), p. 136

7. Bullington, K. and J.M. Fraser, "Engineering Aspects of TASI," Bell System Technical Journal, 38, (March 1959), pp. 353-364.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE *(When Data Entered)*

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER ESD-TR-77-228	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE *(and Subtitle)* Measurements of Voice and Data Queue Behavior in a PVC Network Link		5. TYPE OF REPORT & PERIOD COVERED Technical Note
		6. PERFORMING ORG. REPORT NUMBER Technical Note 1977-37
7. AUTHOR*(s)* Paul Demko		8. CONTRACT OR GRANT NUMBER*(s)* F19628-76-C-0002
9. PERFORMING ORGANIZATION NAME AND ADDRESS Lincoln Laboratory, M.I.T. P.O. Box 73 Lexington, MA 02173		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS Program Element No. 33126K
11. CONTROLLING OFFICE NAME AND ADDRESS Defense Communications Agency 8th Street & So. Courthouse Road Arlington, VA 22204		12. REPORT DATE 25 August 1977
		13. NUMBER OF PAGES 30
14. MONITORING AGENCY NAME & ADDRESS *(if different from Controlling Office)* Electronic Systems Division Hanscom AFB Bedford, MA 01731		15. SECURITY CLASS. *(of this report)* Unclassified
		15a. DECLASSIFICATION DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT *(of this Report)*

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT *(of the abstract entered in Block 20, if different from Report)*

18. SUPPLEMENTARY NOTES

None

19. KEY WORDS *(Continue on reverse side if necessary and identify by block number)*

Packetized Virtual Circuit
voice queue
data queue
multiplexing
vocoding technique
data packets

20. ABSTRACT *(Continue on reverse side if necessary and identify by block number)*

This note describes measurements made on a computer simulation of a model of a Packetized Virtual Circuit (PVC) network link. The simulation models a population of speakers in conversation, and a Poisson data source. Such variables as buffer space requirements, packet loss and delay, and link utilization are investigated as functions of voice and data loads on the system. Initial results indicate that voice and data can be satisfactorily integrated on a 1.544 Mbps packetized communications link with total link utilizations greater than 90%.

DD FORM 1473 1 JAN 73 EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE *(When Data Entered)*